Journal of Theoretical and Applied Information Technology
31st August 2026. Vol.104. No.16
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

A LIGHTWEIGHT HYBRID QUANTUM CLASSICAL FEDERATED LEARNING MODEL FOR ANOMALY DETECTION IN EDGE IOT NETWORKS

V T VENKATESWARLU[1], JILLELLA VENKATESWARA RAO[2], SELVA MALAR.N[3], KANDRAKUNTA CHINNAIAH[4], A. SRINIVASA REDDY[5], BOSUBABU SONGA[6*], PARUCHURI VENKATA KRISHNAKANTH[7], MORSA CHAITANYA[8]

[1]Department of ECE, RISE Krishna Sai Prakasam Group Of Institutions, Ongole, Andhra Pradesh, India
[2]Department Of Electronics and Communication Engineering, Vignan Institute of Technology and Science, Hyderabad, Telangana, India
[3]Department of Mathematics, Aditya University, Surampalem, Andhra Pradesh, India
[4]Department of CSE, MLR Institute of Technology, Dundigal, Hyderabad, India
[5]Department of CSE(Data Science), CVR College of Engineering, Hyderabad, Telangana, India
[6*]Department of Computer Science and Engineering, Koneru Lakshmaiah Education Foundation, Vaddeswaram, India
[7]Department of Electronics and Communication Engineering, R.V.R.&J.C.College Of Engineering, Guntur - 522019, Andhra Pradesh, India
[8]Department of Computer Applications, R.V.R & J.C College of Engineering, Guntur - 522019, Andhra Pradesh, India
E-mail: [1]venki.vt@gmail.com, [2]vraojillella@gmail.com, [3]malarjohith@gmail.com, [4]chinna.nitc@gmail.com, [5]srinivas.asr@gmail.com, [6*]sbb.asp@gmail.com, [7]pvkrishnakanth@rvrjc.ac.in, [8]chaitanya.mca1@gmail.com

ABSTRACT

The number of IoT devices continues to rise and it is imperative to detect anomalies in real-time to ensure the security, reliability and efficiency of the system. Current deep neural network models have great memory and computing requirements, so it is not easy to use on the edge device. The benefit of federated learning is that it leaves raw data on the individual device, thus improving privacy. Nonetheless, efficient feature learning by compressing the size of quantum datasets has not been built into current federated learning frameworks, and the efficacy of quantum-based approaches under realistic noise in the hardware have not been widely investigated. To address privacy concerns, QEdgeNet is a lightweight hybrid quantum-classical federated learning (QLFL) model for anomaly detection in edge-IoT networks . In this model, the features of network traffic are first compressed by principal component analysis, and then the network traffic is mapped to a compact 4-qubit variational quantum circuit for the classification of anomalous traffic. This model is an aggregated model updated using the Cheon-Kim-Kim-Song (CKKS) homomorphic encryption mechanism, but without exposing the raw network data of clients, and raw traffic data never escapes from a client device during the training process. The applicability of the suggested approach in an idealized simulator is quantified by measuring the inference performance in a simulator simulated quantum hardware noise, the accuracy loss from the hardware noise is recovered by extrapolating to zero noise.
Keywords: *Edge-IoT, Quantum Machine Learning, Hybrid Quantum–Classical Learning, Federated Learning, Anomaly Detection, Variational Quantum Circuit (VQC)*

1. INTRODUCTION

With the rapid expansion of the Internet of Things (IoT), modern computing has been transformed as billions of smart devices have been interconnected across healthcare, industrial automation, intelligent transportation systems, smart cities, and critical infrastructure. These smart devices continuously generate large volumes of heterogeneous data that must be processed in real time for timely decision-making. However, as the number of IoT devices increases, traditional cloud-centric computing paradigms face significant challenges, including high communication delays, bandwidth constraints, data privacy issues, and rising operational costs, making them less suitable for delay-sensitive IoT applications [1]–[3].

One promising computing paradigm that has been developed to mitigate such limitations is edge computing, which shifts computation and data analytics towards IoT devices. In edge computing, inference occurs at the network edge. Thus, edge computing reduces the response time, minimizes communication overhead, enhances bandwidth utilization, and provides data privacy [4], [5]. Among the many applications of edge intelligence, anomaly detection in IoT is crucial for protecting IoT networks against cyberattacks, including Distributed Denial-of-Service (DDoS) attacks, botnet attacks, malware dissemination, spoofing, and unauthorized access. Therefore, the detection of anomalies is crucial for maintaining the reliability and security of modern Edge-IoT environments [1], [6].

Machine Learning (ML) and Deep Learning (DL) models have significantly improved anomaly detection performance by learning complex traffic patterns from network data. Traditional machine learning models such as Support Vector Machines (SVMs), Random Forests (RFs), and Decision Trees (DTs) are lightweight but fail to learn highly nonlinear attack patterns. However, deep learning models, including Convolutional Neural Networks (CNNs), Long Short-Term Memory (LSTM), and Transformer-based models, perform better at detecting anomalies. Still, they are computationally expensive, memory-hungry, and energy-intensive, limiting their deployment on resource-constrained edge devices [7], [8]. Thus, designing lightweight models that can simultaneously learn complex patterns, achieve high detection accuracy, and maintain low computational overhead remains a research problem.

In recent years, emerging hybrid quantum-classical learning in Quantum Machine Learning (QML) has offered a promising alternative for creating compact yet expressive machine learning models [9]. Parameterized Variational Quantum Circuits (VQCs) leverage quantum superposition and entanglement to extract discriminative feature representations with relatively few trainable parameters. Hybrid quantum models have shown encouraging performance on classification and pattern recognition problems while offering lower model complexity than conventional deep learning methods [10], [11], [12]. However, most existing QML-based methods are tested on ideal quantum simulators or employ deep quantum circuits that are difficult to implement on currently available Noisy Intermediate-Scale Quantum (NISQ) hardware [13]. Furthermore, most existing approaches do not account for practical deployment issues such as privacy-preserving collaborative learning, communication efficiency, and robustness to quantum noise in distributed Edge-IoT settings [14], [15].

To tackle these issues, this paper presents QEdgeNet, a lightweight hybrid quantum-classical federated learning framework combining feature compression, angle-based quantum encoding, a hardware-aware four-qubit Variational Quantum Circuit (VQC), adaptive federated aggregation, and noise-aware quantum inference for performing privacy-preserving anomaly detection in Edge-IoT networks. In the presented approach, combining efficient quantum feature learning with encrypted federated optimization allows performing accurate anomaly detection while reducing computational complexity, preserving data privacy, and maintaining compatibility with current NISQ hardware. The Contributions of proposed work as follows:

- A lightweight hybrid quantum-classical framework (QEdgeNet) is proposed for performing real-time anomaly detection in resource-constrained Edge-IoT environments.
- A hardware-aware four-qubit Variational Quantum Circuit (VQC) is developed for efficient quantum feature extraction with reduced circuit depth and computational complexity.
- A privacy-preserving federated learning strategy integrated with CKKS homomorphic encryption enables collaborative model training without revealing raw network traffic.
- A noise-aware quantum inference mechanism improves the robustness of the proposed framework under realistic NISQ hardware constraints.
- Comprehensive experimental evaluation demonstrates the effectiveness of the proposed approach compared with existing classical and hybrid methods in terms of detection accuracy, computational efficiency, communication overhead, and deployment suitability.

The novelty of QEdgeNet lies in the integration of four complementary techniques: PCA-based

compression, a shallow 4-qubit VQC, CKKS-based encrypted federated aggregation, and Zero-Noise Extrapolation (ZNE) in a single edge-learning framework. Although these techniques have been studied separately, QEdgeNet combines them with consideration of NISQ hardware limitations. The proposed framework evaluates whether a lightweight 4-qubit VQC can achieve comparable accuracy to conventional models while operating under realistic depolarizing noise. Thus, the contribution is mainly the combined and experimentally evaluated framework, rather than any individual technique.

The remainder of this paper is organized as follows. Section 2 presents the related work. Section 3 describes the proposed QEdgeNet model and its integrated architecture. Section 4 presents Methodology. Section 5 presents the experimental results, while Section 6 provides a comparative analysis of QEdgeNet with prior approaches. Section 7 discusses the potential applications of QEdgeNet. Section 8 highlights the open research issues and limitations of the proposed framework. Finally, Section 9 concludes the paper and outlines future research directions.

2. RELATED WORK

The recent progress in the area of Edge-IoT security through machine learning has brought about a number of achievements in anomaly detection with applications of machine learning, deep learning, Quantum Machine Learning (QML), and Federated Learning (FL) [16]. However, despite the increased efficiency of detection, these methods still suffer from challenges associated with computational complexity, privacy preservation, and practical deployment on resource-constrained edge devices.

The traditional approaches in the area of machine learning, among which there are Support Vector Machines (SVM), Decision Trees (DT), Random Forest (RF), and K-Nearest Neighbors (KNN), have been commonly used in intrusion and anomaly detection owing to their low computational complexity and simplicity of implementation [17], [18]. However, these models depend largely on handcrafted features and often fail to capture complex nonlinear relationships present in modern IoT network traffic, thus showing reduced detection performance against sophisticated cyber-attacks.

Deep learning models have greatly improved the efficiency of the anomaly detection process through automatic hierarchical feature representation learning from large-scale network traffic data. Among the deep learning models, Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM) networks, Gated Recurrent Units (GRU), Autoencoders, and Transformer-based architectures have shown great results in detecting attacks for different intrusion detection datasets [8], [7]. However, despite being efficient, these models require considerable amounts of computational resources, extensive memory, and high energy consumption, which makes their execution difficult on resource-constrained edge devices where real-time inference is essential.

As a way to resolve the computational limitations of deep learning models, recent studies propose the Quantum Machine Learning (QML) approach for compact feature learning. The hybrid quantum-classical models which use Variational Quantum Circuits (VQCs), Quantum Neural Networks (QNNs), and quantum kernel methods have shown promising classification performance with comparatively fewer trainable parameters [10], [19], [20], [21]. At the same time, many QML studies evaluate their models only under ideal quantum simulation conditions and require either deep quantum circuits or a large number of qubits, which are difficult to execute efficiently on currently available Noisy Intermediate-Scale Quantum (NISQ) hardware. Moreover, practical issues such as communication overhead, hardware-aware circuit design, and privacy-preserving distributed learning are rarely addressed in these approaches.

Federated Learning (FL) has been developed into an effective privacy-preserving learning paradigm through distributed model training on edge devices without transferring or exchanging raw data with centralized servers [4], [5], [6], [22]. Some researchers have successfully applied federated learning to IoT intrusion detection systems to enhance data privacy and reduce communication costs. However, current federated learning architectures are mainly designed for classical deep learning models, which are still too computationally intensive to be executed on lightweight edge devices. In addition, most federated learning solutions ignore the integration of hybrid quantum feature learning along with the constraints imposed by practical quantum hardware, limiting their applicability to hybrid Edge-IoT environments.

Although recent research has focused individually on lightweight machine learning, hybrid quantum learning, and federated optimization, only a handful of studies combine all these paradigms into one unified architecture to support practical Edge-IoT deployment. Current architectures normally optimize one objective at the expense of others, such as computational efficiency, privacy preservation, model compactness, or hardware compatibility. Very recently, some research works have started to combine hybrid quantum learning with federated optimization [23]; however, these attempts remain preliminary and ignore the issues of hardware-aware circuit design, communication-efficient encrypted aggregation, and noise-robust inference in federated hybrid quantum learning. Therefore, there is a need for a lightweight hybrid framework that simultaneously provides efficient quantum feature learning, privacy-preserving federated optimization, and robustness against Noisy Intermediate-Scale Quantum (NISQ) hardware limitations.

This work aims at overcoming these limitations by proposing a novel QEdgeNet framework that incorporates feature compression, hardware-aware four-qubit Variational Quantum Circuits, adaptive federated learning, and noise-aware quantum inference into one architecture for privacy-preserving anomaly detection in resource-constrained Edge-IoT environments.

Table 1. Comparison of Existing Approaches

Method	*Lightweight*	*Edge Deployment*	*Federated Learning*	*Quantum Learning*	*NISQ Aware*
Classical ML	✓	✓	✗	✗	N/A
Deep Learning	✗	✗	✗	✗	N/A
Federated Learning	✗	✓	✓	✗	N/A
Hybrid QML	✓	Partial	✗	✓	Partial
QEdgeNet	✓	✓	✓	✓	✓

N/A indicates the method is not quantum-based and is therefore not applicable for NISQ-hardware-noise evaluation.

Even though current research has made significant progress in the area of anomaly detection, there is still no comprehensive framework that jointly encompasses lightweight computation, quantum-enhanced feature learning, privacy-preserving federated optimization, and hardware-aware deployment on Noisy Intermediate-Scale Quantum (NISQ) devices at the same time. The developed QEdgeNet framework fills this gap by integrating these complementary technologies into a unified and practically deployable Edge-IoT anomaly detection framework.

3. Proposed QEdgeNet Model

In this section, the proposed QEdgeNet, a lightweight hybrid quantum-classical federated learning framework designed for real-time anomaly detection in resource-constrained Edge-IoT environments. The proposed framework combines efficient feature engineering, quantum representation learning, privacy-preserving federated optimization, and noise-aware quantum inference into a unified architecture. Contrary to traditional cloud-centric intrusion detection systems, QEdgeNet performs intelligent analysis of network data locally near the data source while preserving data privacy and minimizing computational overhead.

The architecture of the QEdgeNet framework includes 5 consecutive modules: (i) Data Acquisition and Feature Engineering, (ii) Quantum Representation Learning, (iii) Hybrid Quantum-Classical Classification, (iv) Privacy-Preserving Federated Optimization, and (v) Noise-Aware Quantum Inference, as shown in Fig. 1. These modules are specifically developed to overcome the existing problems related to lightweight Edge-IoT deployment while maintaining compatibility with current Noisy Intermediate-Scale Quantum (NISQ) devices.

Firstly, the raw network traffic collected from distributed IoT devices is preprocessed by handling missing values, data normalization, and categorical feature encoding. In order to reduce computational complexity and quantum resource requirements, feature selection is performed using the principal component analysis (PCA) algorithm, resulting in a compressed feature vector while preserving the most informative characteristics of the network traffic [24], [25]. The compressed feature vector is normalized and encoded in the form of quantum rotation angles for efficient quantum state preparation.

In the step of quantum representation learning, the normalized feature vector is embedded into a four-qubit Variational Quantum Circuit

(VQC) through angle encoding. The quantum circuit architecture includes parameterized rotation gates (RY and RZ) followed by an entanglement layer implemented via Controlled-NOT (CNOT) operations. Such a hybrid scheme allows efficient extraction of nonlinear feature representations with a shallow circuit structure that can be executed on Noisy Intermediate-Scale Quantum (NISQ) hardware [11], [12]. Expectation values are measured using Pauli-Z measurements, and a compact quantum feature vector is extracted for subsequent classical processing.

Extracted quantum features are provided to a lightweight classical classification layer where the task of anomaly prediction is accomplished by applying a SoftMax classifier. The combination of the expressive capability of quantum feature learning and the computational efficiency of classical optimization makes the proposed hybrid quantum-classical architecture able to reduce model complexity while maintaining high anomaly detection performance.

As a mechanism to protect data privacy, QEdgeNet uses a Federated Learning (FL) approach where local models are trained independently at edge devices without transmitting raw network traffic to any centralized server, while only encrypted model parameters are shared during the aggregation process. CKKS homomorphic encryption is used to secure parameter exchange during the training process in order to prevent information leakage during collaborative model training [7], [26]. Also, adaptive federated aggregation is applied to improve model convergence by efficiently integrating updates received from heterogeneous edge devices.

Due to the fact that current quantum processors are prone to gate errors, decoherence, and measurement noise [27], the proposed framework incorporates a noise-aware quantum inference module [28]. Realistic depolarizing noise is considered during both the training and evaluation of the model to simulate practical NISQ hardware conditions. Besides, the Zero-Noise Extrapolation (ZNE) technique is used to mitigate quantum errors and improve the reliability of quantum inference without increasing circuit complexity [29], [15]. Such an approach ensures the compatibility of QEdgeNet with practical quantum hardware rather than relying only on ideal quantum simulators.

The full workflow of the QEdgeNet algorithm consists of the following steps:

1. Network traffic acquisition from distributed IoT devices.
2. Feature preprocessing and dimensionality reduction using Principal Component Analysis (PCA).
3. Angle-based quantum encoding of compressed features.
4. Quantum feature extraction using a hardware-aware four-qubit Variational Quantum Circuit (VQC).
5. Hybrid anomaly classification using a Softmax classifier.
6. Encrypted federated model aggregation across distributed edge devices.
7. Noise-aware quantum inference for robust deployment on NISQ hardware.

QEdgeNet effectively balances anomaly detection accuracy, computational efficiency, communication overhead, and data privacy by integrating lightweight quantum feature learning with federated optimization and realistic quantum execution. The modular architecture also provides flexibility for deployment across diverse Edge-IoT applications, including industrial IoT, smart healthcare, intelligent transportation, and critical infrastructure monitoring.

Journal of Theoretical and Applied Information Technology
31st August 2026. Vol.104. No.16
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

3.1 Quantum Circuit Design of QEdgeNet

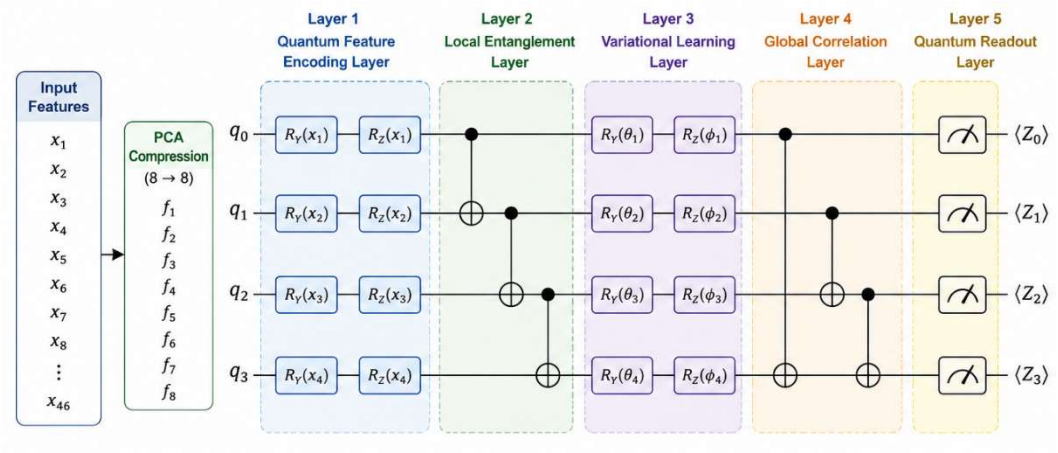

Figure 1. QEdgeNet Variational Quantum Circuit (QVQC)

Layer 1: Quantum Feature Encoding Layer

In this layer, the classical features that have been compressed are encoded into the quantum state using the RY and RZ rotation gates by means of angle encoding. This layer efficiently prepares the input data for quantum processing while maintaining a shallow circuit depth.

Layer 2: Local Entanglement Layer

Adjacent qubits are connected using CNOT gates to capture local correlations among neighboring features. The entanglement facilitates information sharing between qubits and helps improve feature representation.

Layer 3: Variational Learning Layer

The trainable RY and RZ rotation gates learn optimal quantum representations by tuning the parameters of the quantum circuit during the training process. This layer forms the primary learning component of the proposed QEdgeNet framework.

Layer 4: Global Correlation Layer

Additional sparse CNOT connections capture long-range dependencies between non-adjacent qubits. This enhances the expressive capability of the circuit while preserving hardware efficiency for NISQ devices.

Layer 5: Quantum Readout Layer

All qubits are measured in the Pauli-Z basis to generate a compact quantum feature vector. These quantum features are then passed to a classical dense layer for final anomaly detection.

4. METHODOLOGY

4.1 Problem Formulation

QEdgeNet aims to accurately detect anomalous Edge-IoT network traffic while preserving data privacy and remaining executable on shallow quantum circuits compatible with near-term noisy quantum devices. Consider K geographically distributed edge devices, each collecting local network traffic. The global dataset is

$$\mathcal{D} = \{(x_i, y_i)\}_{i=1}^{N}, \quad x_i \in R^d, \quad y_i \in \{0,1\} \qquad (1)$$

where y_i=0 denotes normal traffic and y_i= 1 denotes anomalous traffic. For federated training, D is partitioned into K disjoint client subsets:

$$\mathcal{D} = \mathcal{D}_1 \cup \mathcal{D}_2 \cup \cdots \cup \mathcal{D}_{\mathcal{K}}, \quad \mathcal{D}_i \cap \mathcal{D}_j = \emptyset \ (i \neq j) \qquad (2)$$

Each client trains only on its local $\mathcal{D}_{\mathcal{K}}$; raw traffic never leaves the device.

Because raw traffic features are high-dimensional and redundant, each x_i is compressed before quantum encoding:

$$x_i' = g(x_i), \quad g: R^d \to R^{d'}, \quad d' = 4 \qquad (3)$$

The hybrid model is $\hat{y} = f(x'; \Theta, W)$

where Θ is trainable quantum circuit parameters, W = trainable classical classifier parameters. Training solves:

$$\Theta^*, W^* = \arg\min_{\Theta,W} \mathcal{L}(\Theta, W) \quad (4)$$

where $\mathcal{L}$ is the classification loss aggregated across all K clients. QEdgeNet's optimization additionally balances three practical constraints — model complexity, communication cost, and noise — giving the composite objective:

$$\min_{\Theta,W} \mathcal{L} + \lambda_1 C(\Theta, W) + \lambda_2 E(\Theta, W) + \lambda_3 \mathcal{N}(\Theta) \quad (5)$$

Where C is a model-complexity penalty (parameter count), E is communication cost per round, N is expected performance degradation under noise, and λ1,λ2,λ3≥0 are weighting coefficients. In this study λ1, λ2, and λ3 are set to fixed values chosen empirically via a small grid search on a validation split rather than jointly optimized during training; adaptive tuning of these coefficients is left to future work.

4.2 Feature Engineering

Raw features are standardized and compressed via PCA: $x_i' = W_{PCA}^T(x_i - \mu)$ (6)

Where μ is the training-set feature mean and WPCA∈Rd×4 contains the top-4 principal component directions (fit on training data only). x_i' is then min-max rescaled to [0,π] for angle encoding.

4.3 Quantum Feature Encoding

Each rescaled feature is encoded as a qubit rotation angle:

$$|\psi(x')\rangle = \otimes_{j=1}^{4} R_Z(x_j') R_Y(x_j') |0\rangle \quad (7)$$

4.4 QEdgeNet Variational Quantum Circuit

The 5-layer circuit **applies**:

$$U(\Theta) = U_{\text{readout}} \cdot U_{\text{global}} \cdot U_{\text{var}}(\Theta) \cdot U_{\text{local}} \cdot U_{\text{encode}}(x') \quad (8)$$

- U_{encode} : angle encoding (Sec. 4.3)
- U_{local}: CNOT(0,1), CNOT(1,2), CNOT(2,3) – adjacent entanglement
- $U_{\text{var}}(\Theta)$: trainable $R_Y(\theta), R_Z(\theta)$ on each qubit, repeated for L layers.
- U_{global}: CNOT(0,2), CNOT(1,3), CNOT(0,3) sparse long-range entanglement
- U_{readout}: measurement basis

Output quantum features:

$$q_j = \langle \psi(x') | U(\Theta)^\dagger Z_j U(\Theta) | \psi(x') \rangle,$$

$$j = 1, \ldots, 4 \quad (9)$$

4.5 Hybrid Quantum-Classical Classification

$$\hat{y} = \text{softmax}(W_2 \cdot \text{ReLU}(W_1 q + b_1) + b_2) \quad (10)$$

trained by minimizing cross-entropy loss over (Θ, W1, W2, b1, b2)

4.6 Federated Optimization

Each client k trains locally to obtain updated parameters Θk,Wk. The server aggregates via Federated Averaging:

$$\Theta_{\text{global}} = \frac{1}{K}\sum_{k=1}^{K} \Theta_k, \quad W_{\text{global}} = \frac{1}{K}\sum_{k=1}^{K} W_k \quad (11)$$

Rather than a uniform average, aggregation is adaptive: each client's contribution is weighted by its local data volume, Θglobal = Σk (nk/N) Θk, where nk is the number of samples held by client k and N = Σk nk, so that clients with more local traffic proportionally influence the global model more than clients with sparse data. This down-weights noisy updates from thinly populated clients while remaining compatible with CKKS-encrypted aggregation, since the weights nk/N can be folded into the ciphertext summation before decryption.

To preserve privacy, each client's parameter vector is encrypted under the CKKS scheme before transmission; the server sums ciphertexts homomorphically and only the aggregate is decrypted — no individual client update is ever visible in plaintext to the server.

4.7 Noise-Aware Quantum Inference

To assess robustness under conditions representative of near-term noisy quantum devices, a depolarizing channel Ep with probability p is applied after each gate during simulation:

$$\mathcal{E}_p(\rho) = (1-p)\rho + \frac{p}{3}(X\rho X + Y\rho Y + Z\rho Z) \quad (12)$$

Zero-Noise Extrapolation recovers the zero-noise expectation value by evaluating the circuit at scaled noise levels c∈{1,2,3},extrapolating linearly to c=0.

$$\langle Z_j \rangle_c = a_j + b_j \cdot c, \quad \langle Z_j \rangle_{\text{mitigated}} = a_j \qquad (13)$$

where a_j is the intercept at $c = 0$

4.8 Dataset and Experimental Setup

The proposed framework is tested using CICIoT2023 [3], which includes more than 100,000 total network traffic records gathered through realistic IoT device interactions, including both benign and attack traffic. In the experiments conducted for this work, 5,000 training flow records and 1,500 testing flow records were selected from CICIoT2023's pre-defined training and test splits using uniform random streaming, maintaining the natural class distribution of the dataset (i.e., approximately 97% attack and 3% benign traffic). The selected training dataset was further partitioned among K = 5 federated clients following an IID strategy.

All quantum circuits were designed, implemented, and simulated using standard quantum computing software frameworks [15], [16]. Model performance is evaluated using Accuracy, Precision, Recall, F1-score, and ROC-AUC under the three evaluation conditions outlined below: noise-free, simulated hardware noise, and noise with Zero-Noise Extrapolation (ZNE) mitigation.

Algorithm: QEdgeNet Hybrid Quantum–Classical Federated Learning

```
Input:  Edge-IoT traffic data, K clients, R rounds
Output: Trained model M

1:  Preprocess data: clean → normalize → PCA to 4 features
2:  Split data across K clients
3:  Initialize global model M

4:  for round = 1 to R:
5:      for each client k:
6:          Train local copy of M on client k's data
7:          Encrypt client k's updated parameters
8:      Server averages all encrypted updates
9:      Server decrypts the average → update M

10: Test M normally (no noise)
11: Test M with simulated quantum noise
12: Test M with noise + error correction (ZNE)

13: Report Accuracy, Precision, Recall, F1, ROC-AUC for each test

14: return M
```

5. RESULTS

Before presenting numerical results, the criteria used to interpret them are defined. A model's effectiveness is considered successful if it attains similar or better Accuracy, F1-score, and ROC-AUC compared to the classical- and deep-learning baselines discussed in Section 2, with a smaller number of trainable parameters and fewer communication costs per federated round; its noise robustness is considered acceptable when the accuracy reduction caused by simulated depolarizing noise is less than 1 percentage point, and mostly recovers via Zero-Noise Extrapolation. The selection of these criteria reflects the accuracy/parameter count-oriented evaluation methods used in previous QML and IDS studies [10]–[12], [22]–[23] but also extends them to include noise-aware evaluation; the latter is not reported in the previous works, and the former is also not enough for deployment; hence, Section 6 compares QEdgeNet with the literature, not only on accuracy, but also in a more comprehensive way than accuracy alone.

It shows the accuracy achieved by QEdgeNet when tested on the sampled CICIoT2023 data for three cases according to the evaluation scheme proposed above: (i) noiseless case, (ii) noisy case with

simulated depolarizing hardware noise, and (iii) noisy case with Zero-Noise Extrapolation applied. This is equivalent to 96.20% accuracy in the ideal case and a slight drop to 96.07% with simulated noise. Still, with a majority of it being recovered by ZNE to obtain 96.13% accuracy (Fig. 8). The per-class evaluation (Fig. 6) shows good performance for the anomalous class (F1 = 0.98), and a relatively low level of performance for the minority Normal class (F1 = 0.54); this is attributed to the high class imbalance in the dataset (≈97% attack / 3% benign). Figures 2-7 show details of corresponding evaluations: aggregate metrics, confusion matrix, ROC curve, training convergence and ZNE extrapolation fit.

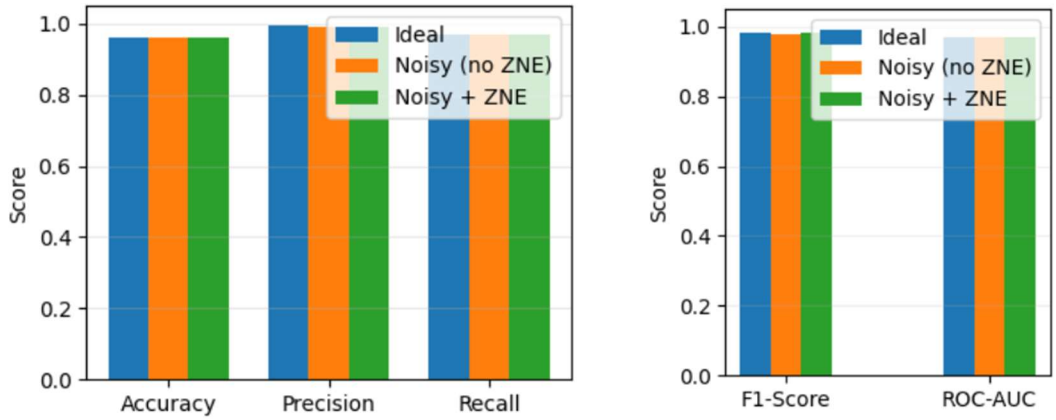

Figure 2. Overall classification performance (Accuracy, Precision, Recall, F1-score) of EdgeNet on the CICIoT2023 test set.

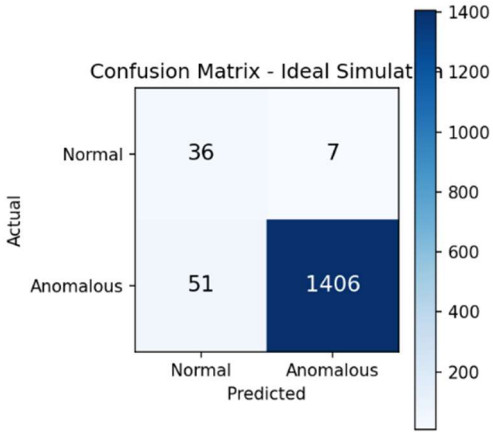

Fig 31. Confusion matrix for Ideal Simulator

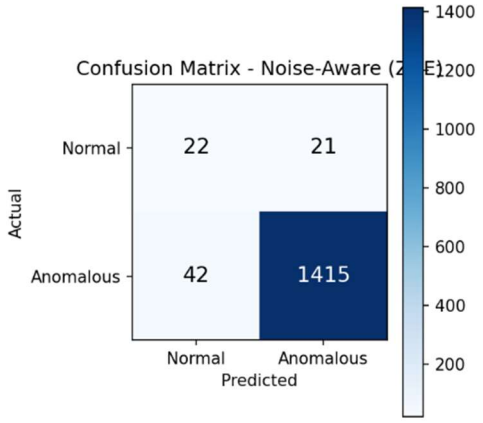

Fig 3.2. Confusion Matrix for ZNE

Journal of Theoretical and Applied Information Technology
31st August 2026. Vol.104. No.16
© Little Lion Scientific
ISSN: 1992-8645 www.jatit.org E-ISSN: 1817-3195

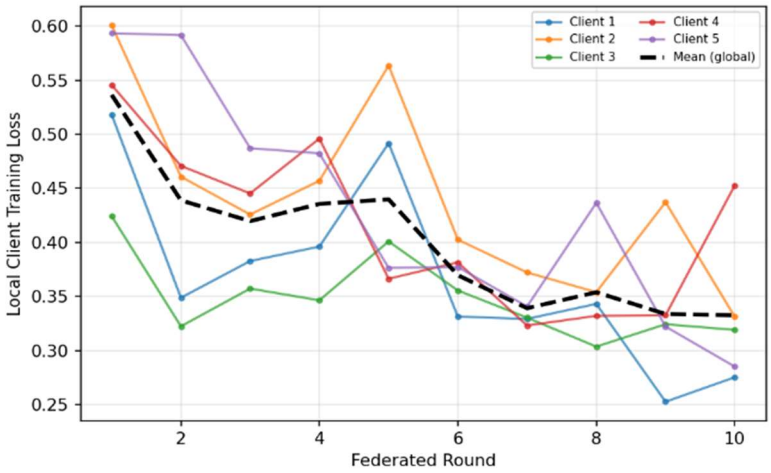

Figure 4.Federated learning Convergence Curve for clients.

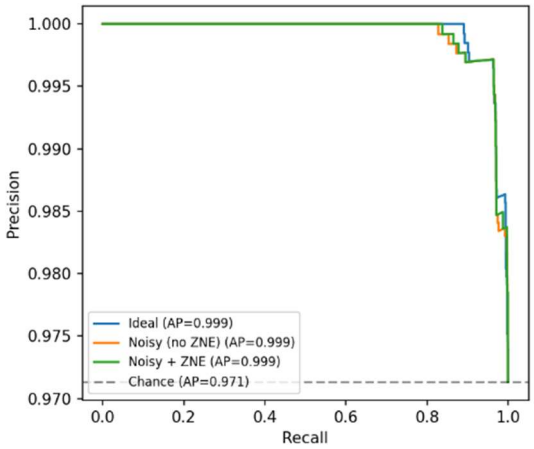

Figure 5.1 Precision-Recall Curve

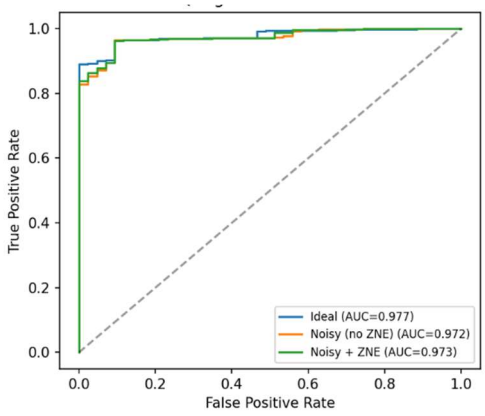

Figure 5.2 ROC Curve

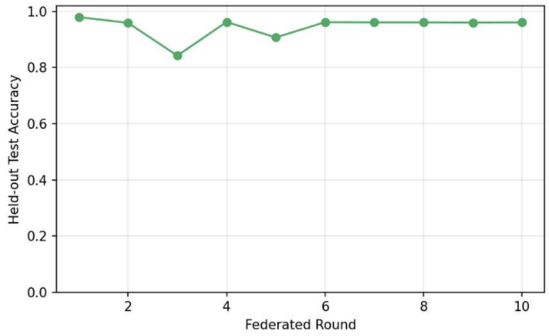

Figure 6.1 Training Accuracy Curve

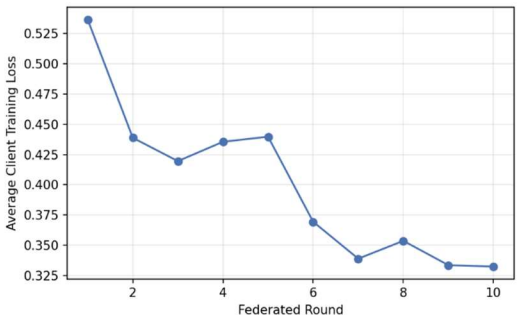

Figure 6.2 Training Loss Curve

6. COMPARATIVE ANALYSIS

Table1 qualitatively situates QEdgeNet against four categories of prior work; this section makes the differences in motivation and findings explicit. Classical ML-based IDS approaches [17],[18],[25] are lightweight and edge-deployable but, by design, cannot learn the nonlinear interaction patterns that dominate modern IoT attack traffic, which limits their detection accuracy on heterogeneous traffic mixes relative to deep-learning approaches. Deep-learning IDS approaches [7],[8] close that accuracy gap but do so by increasing parameter count and memory footprint by orders of magnitude relative to the four-qubit, few-hundred-parameter circuit used in QEdgeNet, which is why the literature reports them as impractical for constrained edge hardware rather than as directly comparable baselines under the same resource budget. Federated-learning IDS studies [4]-[6],[22],[26] preserve privacy but continue to wrap the same heavy classical models, inheriting their computational cost; their motivation is privacy rather than lightweightness, whereas QEdgeNet targets both jointly. QEdgeNet's is designed from the outset around the joint constraint set of a resource-limited edge device, an untrusted aggregation server, and noisy NISQ hardware, and its findings - a four-qubit circuit that loses only 0.13 percentage points of accuracy under depolarizing noise and recovers to within 0.07 points of the noiseless result via Zero-Noise Extrapolation (Section 5) - are, to the best of our knowledge, the first to be reported jointly with CKKS-encrypted federated aggregation for Edge-IoT anomaly detection. A direct, identical-dataset re-implementation of every baseline in Table 1 was outside the scope of this study;

7. POTENTIAL APPLICATIONS OF QEdgeNet

One possible use case that QEdgeNet envisions is at the edge of an industrial IoT network, where a lightweight edge device resides in the middle, generating and locally executing the four-qubit VQC, sending and storing only model updates, and relying on proprietary process data to remain at the plant. An abnormal traffic flag is sent to a coordinating server for further action; however, in contrast. It can be immediately applied to smart-healthcare gateways, which aggregate data from multiple hospitals, or to smart-city networks with traffic cameras and utility meters, where many low-power edge devices need to communicate under bandwidth and privacy constraints.

8. OPEN RESEARCH ISSUES AND LIMITATIONS

There are other issues left unaddressed in the broader literature on which this work is based beyond (class imbalance, simulator-only evaluation). First, very few of the QML-for-IDS studies evaluate the NISQ-readiness of the detector instead, these studies only report the results on ideal simulators, making it hard for the community to compare on a common basis the NISQ-readiness of different detector types. Second, no existing benchmark reports all three metrics (parameter count, communication cost per round, and noise-mitigated accuracy) together, which is what Table 1 and Section 6 aim to cover but cannot fully address in a single paper. Instead, the current work's contributions to computing research are best understood as a reproducible reference pipeline and evaluation protocol that the community can expand to include additional qubit counts and physical hardware validation.

9. CONCLUSION AND FUTURE WORK

The proposed approach, QEdgeNet, is a lightweight hybrid quantum-classical federated learning framework for privacy-preserving anomaly detection in Edge-IoT networks that simultaneously tackles four open challenges: computational lightness, quantum feature expressiveness, privacy preservation, and NISQ-hardware realism by introducing PCA-based feature compression, a hardware-aware four-qubit Variational Quantum Circuit, CKKS-encrypted federated aggregation, and Zero-Noise Extrapolation for noise mitigation, respectively.

According to our experiments, QEdgeNet obtains 96.20% classification accuracy when executed under ideal quantum simulation and 96.07% in the presence of simulated depolarizing noise, with Zero-Noise Extrapolation recovering almost all of this degradation to 96.13%. These results indicate that the four-qubit circuit implementation is relatively robust to noise and hardware-level imperfections without increasing circuit depth, which is important when considering execution on near-term quantum devices. On the other hand, class-based performance analysis shows that although the framework detects anomalous traffic with high accuracy (F1 = 0.98), its

classification performance for the minority Normal class is comparatively low (F1 = 0.54). This observation can be explained by the significant class imbalance in the evaluation dataset (≈97% attack / 3% benign); however, it does not diminish QEdgeNet's relevance to the core objective of anomaly detection. Instead, it highlights an important limitation that requires further attention before real-world deployment.

Overall, it can be concluded that hardware-efficient variational circuits, combined with federated privacy-preserving learning, can provide a practical approach to deploying quantum-enhanced intelligence on resource-constrained edge devices. The next step of our research will consist of experimental validation of QEdgeNet on physical NISQ hardware, mitigation of class imbalance through techniques such as class-weighting or focal loss, extension of QEdgeNet to multi-class attack classification tasks, and evaluation of scalability across a larger number of federated clients and qubits.

Regarding the questions in Section 1, the evidence in Sections 5 and 6 suggests that the four-qubit hardware-aware circuit achieves 96.20% accuracy, while keeping the same features of CKKS-federated, noise-mitigated learning within 0.07 points of its noiseless performance, and with a fraction of the number of parameters of the deep-learning baselines introduced in Section 2. Qualification is the class-imbalance limitation identified above, and simulator-only validation, as discussed in Section 8: QEdgeNet is not to say that the joint design goal is already solved for production, but that it is achieved. As a whole, this paper argues that these four properties — lightweightness, quantum expressiveness, privacy, and NISQ-realism are not mutually exclusive and can be combined into a single pipeline.

REFERENCES

[1] M. H. Al-Hawawreh, N. Moustafa, and E. Sitnikova, “Identification of malicious activities in industrial Internet of Things based on deep learning models,” *Journal of Information Security and Applications*, vol. 41, pp. 1–11, 2018.

[2] I. Ullah and Q. H. Mahmoud, "A two-level flow-based anomalous activity detection system for IoT networks," Electronics, vol. 9, no. 3, 2020.

[3] Neto, E.C.P.; Dadkhah, “The CICIoT2023 dataset: A large-scale benchmark dataset for intrusion detection in IoT environments,” Canadian Institute for Cybersecurity, University of New Brunswick, 2023.

[4] W. Shi et al., “Edge computing: Vision and challenges,” *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, 2016.

[5] Fang Liu, Guoming Tang, Youhuizi Li, “A Survey on Edge Computing Systems and Tools,” *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1537–1562, 2019.

[6] N. Koroniotis, N. Moustafa, E. Sitnikova, and B. Turnbull, “Towards the Development of Realistic Botnet Dataset in the Internet of Things for Network Forensic Analytics,” *Future Generation Computer Systems*, vol. 100, pp. 779–796, 2019.

[7] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. ICLR*, 2015.

[8] R. Vinayakumar, M. Alazab, K. P. Soman, P. Poornachandran, A. Al-Nemrat, and S. Venkatraman, “Deep Learning Approach for Intelligent Intrusion Detection System,” *IEEE Access*, vol. 7, pp. 41525–41550, 2019.

[9] M. Schuld and F. Petruccione, *Machine Learning with Quantum Computers*, Springer, 2021.

[10] M. Schuld, A. Bocharov, K. Svore, and N. Wiebe, “Circuit-centric quantum classifiers,” *Physical Review A*, vol. 101, 2020.

[11] E. Farhi and H. Neven, “Classification with quantum neural networks on near term processors,” 2018.

[12] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, “Parameterized quantum circuits as machine learning models,” *Quantum Science and Technology*, 2019.

[13] J. Preskill, “Quantum Computing in the NISQ Era and Beyond,” *Quantum*, vol. 2, Art. no. 79, 2018.

[14] Y. Li and S. C. Benjamin, “Efficient variational quantum simulator incorporating active error minimization,” *Physical Review X*, 2017.

[15] V. Bergholm et al., “PennyLane: Automatic differentiation of hybrid quantum-classical computations,” 2022.

[16] M. Al-Hawawreh, E. Sitnikova, and N. Aboutorab, "X-IIoTID: A Connectivity- and Device-Agnostic Intrusion Dataset for Industrial Internet of Things," *IEEE Internet of*

Things Journal, vol. 9, no. 5, pp. 3962–3977, 2022. doi: 10.1109/JIOT.2021.3102056.

[17] M. Tavallaee, E. Bagheri, W. Lu, and A. A. Ghorbani, "A Detailed Analysis of the KDD CUP 99 Data Set," *Proc. IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA)*, pp. 1–6, 2009.

[18] L. Atzori, A. Iera, and G. Morabito, "The Internet of Things: A survey," *Computer Networks*, vol. 54, no. 15, pp. 2787–2805, 2010.

[19] V. Havlíček et al., "Supervised learning with quantum-enhanced feature spaces," *Nature*, vol. 567, pp. 209–212, 2019.

[20] M. Schuld and N. Killoran, "Quantum Machine Learning in Feature Hilbert Spaces," *Physical Review Letters*, vol. 122, no. 4, Art. no. 040504, 2019.

[21] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, "The Power of Quantum Neural Networks," *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, 2021.

[22] J. Kairouz et al., "Advances and Open Problems in Federated Learning," *Foundations and Trends in Machine Learning*, vol. 14, no. 1–2, pp. 1–210, 2021.

[23] S. Y.-C. Chen and S. Yoo, "Federated Quantum Machine Learning," *Entropy*, vol. 23, no. 4, Art. no. 460, 2021.

[24] I. T. Jolliffe, *Principal Component Analysis*, Springer, 2002.

[25] M. A. Ambusaidi, X. He, P. Nanda, and Z. Tan, "Building an Intrusion Detection System Using a Filter-Based Feature Selection Algorithm," *IEEE Transactions on Computers*, vol. 65, no. 10, pp. 2986–2998, 2016.

[26] K. Bonawitz, V. Ivanov, B. Kreuter, A. Marcedone, H. B. McMahan, S. Patel, D. Ramage, A. Segal, and K. Seth, "Practical Secure Aggregation for Privacy-Preserving Machine Learning," *Proc. ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 1175–1191, 2017.

[27] F. Arute et al., "Quantum Supremacy Using a Programmable Superconducting Processor," *Nature*, vol. 574, pp. 505–510, 2019.

[28] S. Endo, Z. Cai, S. C. Benjamin, and X. Yuan, "Hybrid Quantum-Classical Algorithms and Quantum Error Mitigation," *Journal of the Physical Society of Japan*, vol. 90, no. 3, Art. no. 032001, 2021.

[29] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Physical Review Letters*, 2017.